MNRAS **000**, 000–000 (0000) Preprint 19 April 2017 Compiled using MNRAS LATEX style file v3.0

Characterising dark matter haloes with computer vision

Julian Merten [1] ⋆, Quim Llorens [2,1] and Hans Winther [3]

[1] *Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, U.K.*
[2] *Facultat de Física, Universitat de Barcelona, Martí i Franquès, 1. E-08028 Barcelona, Catalonia, Spain*
[3] *Institute of Cosmology & Gravitation, University of Portsmouth, Portsmouth, Hampshire, PO1 3FX, U.K.*

arXiv:1704.05072v1 [astro-ph.CO] 17 Apr 2017

© 2017 University of Oxford. All rights reserved.
Submitted to the Monthly Notices of the Astronomical Society.

ABSTRACT

This work explores the ability of computer vision algorithms to characterise dark matter haloes formed in different models of structure formation. We produce surface mass density maps of the most massive haloes in a suite of eight numerical simulations, all based on the same initial conditions, but implementing different models of gravity. This suite includes a standard ΛCDM model, two variations of $f(R)$-gravity, two variations of Symmetron gravity and three Dvali, Gabadadze and Porrati (DGP) models. We use the publicly available `WND-CHARM` algorithm to extract 2919 image features from either the raw pixel intensities of the maps, or from a variety of image transformations including Fourier, Wavelet, Chebyshev and Edge transformations. After discarding the most degenerate models, we achieve more than 60% single-image classification success rate in distinguishing the four different models of gravity while using a simple weighted neighbour distance (WND) to define our classification metric. This number can be increased to more than 70% if additional information, such as a rough estimate of the halo mass, is included. We find that the classification success steeply declines when the noise level in the images is increased, but that this trend can be largely reduced by smoothing the noisy data. We find Zernike moments of the Fourier transformation of either the raw image or its Wavelet transformation to be the most descriptive feature, followed by the Gini coefficient of several transformations and the Haralick and Tamura textures of the raw pixel data eventually pre-processed by an Edge transformation. The proposed methodology is general and does not only apply to the characterisation of modified gravity models, but can be used to classify any set of models which show variations in the 2D morphology of their respective structure.

Key words: dark matter – large-scale structure of Universe – gravitation – methods: numerical.

1 INTRODUCTION

The widely accepted standard model of cosmology is based on general relativity, a cold dark matter component and a cosmological constant. While this ΛCDM model is very successful in describing observations of different kinds on very large scales (Betoule et al. 2014; Anderson et al. 2014; Planck Collaboration et al. 2016a; Hildebrandt et al. 2017), tensions between theoretical predictions and observations arise when entering the realm of non-linear structure formation, at the nexus of cosmology and astrophysics. These problems include the counted number of galaxy clusters as a function of mass and redshift (Planck Collaboration et al. 2016b), the exact shape of the dark matter density profile on scales ranging from galaxy clusters to dwarf galaxies (Newman et al. 2013; Read et al. 2016; Walker & Peñarrubia 2011), the scale and distribution of halo substructure (Kauffmann et al. 1993; Boylan-Kolchin et al. 2012; Schwinn et al. 2017), the diversity in simulated dark matter profiles (Oman et al. 2015) and the unexpected correlation between baryonic and dark matter components of galaxies (McGaugh et al. 2016). A thorough understanding of these tensions is difficult due to the non-linear nature of structure formation. In this regime, theoretical predictions generally rely on the results of cosmological N-body and hydrodynamical simulations.

Several models of structure formation have been proposed to remedy some or all of the aforementioned tensions between observations and the simplest ΛCDM simulations.

⋆ julian.merten@physics.ox.ac.uk

© 0000 The Authors

Such models include, but are not limited to, a more detailed modelling of baryonic effects including feedback processes on cosmological scales via sub-grid physics (among others, Vogelsberger et al. 2014; Dubois et al. 2014; Schaye et al. 2015; McCarthy et al. 2017) and in greater detail with smaller boxes using pristine resolution (Hopkins et al. 2014; Kimm et al. 2015; Nelson et al. 2016); alternative models of gravity (e.g. Clifton et al. 2012, and references therein); and more general models of dark matter, e.g. warm dark matter (e.g. Lovell et al. 2014; Bozek et al. 2016, and references therein) or self-interacting dark matter (Rocha et al. 2013; Peter et al. 2013; Vogelsberger et al. 2014).

This recent emergence of a larger variety of structure formation models, poses the immediate challenge of finding their signatures in astrophysical data. In this work, we aim for characterising the distribution of matter in a very agnostic fashion. While there are many possible descriptors to characterise structure in simulations or real data, our approach is driven by the fact that gravitational lensing (see Bartelmann 2010, for a review) has recently made great progress in delivering a detailed and robust 2D picture of the matter distribution in observed massive haloes (e.g. Johnson et al. 2014; Merten et al. 2015; Jauzac et al. 2015; Massey et al. 2015; Treu et al. 2016; Caminha et al. 2016; Diego et al. 2016, for recent analyses). In general, the characterisation of an observed halo is reduced to a single number such as its mass, or to a two or three parameter density profile under the assumption of spherical symmetry. But the increased number of background galaxies for weak lensing studies (e.g. Clowe et al. 2006; Massey et al. 2007; Merten et al. 2011; Dawson et al. 2012; Jee et al. 2014; Umetsu et al. 2014; Melchior et al. 2015; Harvey et al. 2015; Medezinski et al. 2016) combined with more sophisticated modelling techniques, which combine multi-scale tracers such as weak -and strong lensing (e.g. Bradač et al. 2009; Jullo & Kneib 2009; Merten et al. 2009; Diego et al. 2015; Merten 2016), allows for a detailed reconstruction of the total matter distribution on a range of scales. The full morphological information in such 2D maps shall be used to draw conclusion on the mechanisms of structure formation. This approach avoids the usual compression of valuable information into very few numbers and hence stands a better chance of being able to distinguish the potentially subtle differences between structure formation models.

The scope of this work is a pilot study of the mentioned approach in the context of modified gravity models, but the framework is general and can be used to characterise haloes in any model of structure formation. Our approach uses computer vision techniques to extract the morphological information in surface mass density maps of dark matter haloes. This approach developed from digital image processing (Gonzalez & Woods 2007) and is now routinely used in scientific applications such as medical imaging, microscopy or material sciences. The article is structured as followed: In section 2 we describe our modified gravity simulations, explain how we create surface mass density maps of haloes and introduce our computer vision based image characterisation algorithm together with a simple classifier. We demonstrate our ability to distinguish the different modified gravity models based on this approach in section 3 and discuss these results in section 4. We conclude in section 5 and give an outlook on future applications and improvements of the method.

2 METHODOLOGY

In the following we describe the different steps necessary to use computer vision based image characterisation to classify haloes into underlying structure formation models. Our main data set is a suite of simulations, carried out in eight different models of gravity and using the same initial conditions. From each simulation we choose the most massive systems and produce 2D maps of their projected matter distribution. These maps are smoothed and converted into a standard computer image format. A computer vision algorithm then extracts almost 3000 features from the image into a feature vector. All feature vectors per simulation class form a training set, which is used to classify a test set based on a simple distance metric in image feature space.

2.1 Numerical simulations

Our suite of numerical simulations is described in Winther et al. (2015) and has been used for scientific analyses in Mead et al. (2016), L'Huillier et al. (2017) and von Braun-Bates et al. (2017). All simulations were carried out with a modified version (Llinares et al. 2014) of the adaptive mesh refinement code Ramses (Teyssier 2002) and implement four different models of gravity. Common to all runs is the box size 250 Mpc/h, the number of 512^3 particles and the evolution of the cosmological background which is assumed to follow a linear ΛCDM model with $(\Omega_m, \Omega_\Lambda, h_{100}) = (0.267, 0.733, 0.704)$. All simulation runs were started with identical initial conditions, a most crucial requirement for our study.

Since all the simulations we consider here have the same background evolution, all the differences between the modified gravity simulations and ΛCDM stem from the presence of a fifth-force in the former simulations. The modified gravity models also have what is known as a screening mechanism which is a way of dynamically 'hiding' the modifications of gravity in high density regions (relative to the cosmic mean). This allows these models to pass the stringent constraints from tests of gravity in the Solar System and at the same time giving rise to large modifications on cosmological scales. One can think of these models as effectively introducing a modified gravitational law of the form $F = \frac{GMm}{r^2}(1 + C \cdot \epsilon \cdot e^{-r/\lambda})$. Here C describes the strength of the fifth-force relative to standard gravity, λ is the interaction range of this force and ϵ is a screening factor. The screening factor is 1 if the masses are small and/or are located in a low density environment. Otherwise the screening factor becomes $\epsilon \ll 1$. What constitutes small and large mass haloes in this context depends on the model parameters.

The first gravity model we consider is following standard Newtonian gravity ($C = 0$) and we will dub this model as *lcdm* in the following. The second class are Hu-Sawicki $f(R)$ models (Hu & Sawicki 2007), where the critical mass for a halo to be screened or not is $M_{\rm crit} \sim 10^{13}(f_{R0}/10^{-6})^{1.5} M_\odot$. In this study we look into two of these models, the first one with $|f_{R0}| = 10^{-5}$ and another

one with $|f_{R0}| = 10^{-6}$. We label the two models as *f5* and *f6*, respectively. The coupling strength of these models is $C = \frac{1}{3}$ and they have an interaction range at the present time of 8 Mpc/h (2 Mpc/h) for *f5* (*f6*) which becomes smaller and smaller the further back in time we go. Hence, we expect stronger deviations from *lcdm* in the *f5* model compared to *f6*. We also consider two different manifestations of Symmetron gravity (Hinterbichler & Khoury 2010), labelled *symmA* and *symmB*, with a coupling $C = \sqrt{1-(a_{\rm SSB}/a)^3}$ for $a < a_{\rm SSB}$ and $C = 0$ otherwise. The interaction range is $\lambda = 1$ Mpc/h. The difference between the models is the scale factor of the symmetry breaking which is $a_{\rm SSB} = 0.5$ for *symmA* and $a_{\rm SSB} = 0.33$ for *symmB*. The modifications of gravity do not start to take effect before the scale factor is greater than $a = a_{\rm SSB}$ so we expect larger effects for smaller values of $a_{\rm SSB}$. The final class of models are the normal-branch Dvali–Gabadadze–Porrati (DGP) gravity models (Dvali et al. 2000). The DGP models have an infinite interaction range (just like for normal gravity) and $C \approx \frac{1}{3(1+r_c H(a))}$. The relevant parameter here is the cross-over scale r_c above which the 5-dimensional nature of the theory becomes relevant. Here we study three such scales with $r_c H_0 = 0.5$, $r_c H_0 = 1.2$ and $r_c H_0 = 5.6$, labelled *dgp05*, *dgp12* and *dgp56*, respectively. We expect the larger changes in comparison to the *lcdm* baseline model for smaller values of $r_c H_0$. The final difference between the models is in how the screening, the ϵ factor, behaves. For the Symmetron and $f(R)$ models it is determined by the local value of the gravitational potential Φ and in the DGP models it is determined by the local matter density. For the concrete implementation of the changes in the underlying equations of gravity the interested reader can refer to Llinares et al. (2014) and Winther et al. (2015). For a more thorough discussion about how these models work and modified gravity in general see Clifton et al. (2012) and references therein.

For all eight simulation runs we have snapshots at three different redshifts $z = 0.0$, $z = 0.5$ and $z = 1.0$. For each box and at all redshifts a halo catalogue was produced by the Rockstar halo finder (Behroozi et al. 2013). From each of the eight halo catalogues at three different redshifts we choose the 200 most massive objects, which are not a subhalo of a more massive parent halo. At $z = 0.0$ this results in a typical mass range $1.0 - 7.6 \times 10^{14} M_\odot/h$, at $z = 0.5$ the 200 systems lie between $0.7 - 4.6 \times 10^{14} M_\odot/h$ and for the highest redshift $z = 1.0$ between $0.4 - 3.4 \times 10^{14} M_\odot/h$.

2.2 Image creation

From the particle positions and masses within each simulation snapshot we create surface mass density maps by centring the coordinate system on a user-specified origin and by rotating the coordinate frame to match an, again user-specified, line-of-sight. We then define a cube of given side length around the current origin and with its z-axis aligned with the current line-of-sight. The surface mass density map is created by diving the x-y face of the cube, which is perpendicular to the line-of-sight, in $N_x \times N_y$ cells and by projecting all particles in the cube that lie within the cell while moving along the line-of-sight. From the surface mass density maps we create 8-bit TIFF[1] images by normalising the surface mass density values of the map with respect to a $[0, 255]$ interval.

2.3 Training and test sets

We create images of the 200 most massive haloes in each of the eight gravity models and for all three redshifts. We use a cube size of 10 Mpc/h and we sample the surface mass density with 256×256 cells. We will refer to these cells as pixels from now on. Since we can conveniently exploit the fact that we are working with projected quantities, we use a total of 20 randomly chosen lines-of-sight per halo and hence create for each gravity model and at all redshifts a training set class of 4000 images. In order to test the effects of smoothing on the images we produce six different realisations of each training set, the original un-smoothed version and versions with a Gaussian smoothing with standard deviations of 1, 2, 3, 4 and 5 pixel(s), respectively.

We also create test sets of images to check the performance of the classifier later on. For these, we again use the 200 most massive haloes, for each gravity model and at all redshifts, but along single a line-of-sight which differs from the ones used for the training sets. In order to include a simple noise model we apply Gaussian white noise to each cell of the surface mass density map with signal to noise ratios (SNR) in each pixel of ∞ (no noise applied), 20, 10, 5, 2 and 1 respectively. To all test set images, for all models, redshifts and shot noise levels, we again produce six smoothed realisations using the same Gaussian filters as for the training set.

In summary, we create a training set with eight classes, each class containing 4000 images. For each training set we create six versions with increasing levels of smoothing applied. We create the same number of test set classes, with the difference that they are based on only 200 images per class and that they have different levels of noise applied. These images, split into a test and a training set, containing eight classes each and available at three different redshifts, are the starting point for the subsequent characterisation procedure.

2.4 Characterisation

We use the publicly available software `WND-CHARM`[2] to characterise the halo images. This software was originally designed for medical and biological applications, especially the classification of objects seen under a microscope (Shamir et al. 2008). Many other such algorithms or complete software packages exist (e.g. Bengtsson & Rodenacker 2003; Heller & Ghahramani 2006; Yavlinsky et al. 2006, for a variety of different approaches), but a unique feature of `WND-CHARM` is the fact that it derives a large amount of image descriptors of very different kinds. It is thereby not limited to a specific set of image features, such as e.g. image textures, and hence renders it ideal for our aim in this pilot study: a fully agnostic view on the classification problem using computer

[1] http://www.fileformat.info/format/tiff/index.dir
[2] https://github.com/wnd-charm/wnd-charm

vision with very little a-priori knowledge on what the most descriptive image features will be.

The monochrome, sometimes called grey-scale, 256×256 pixel 8-bit TIFF images are converted by `WND-CHARM` into an image feature vector. While doing so it does not only operate on the raw pixel intensity data, but also on several transformations of the image including its Fourier (F), Chebyshev (C) and wavelet (W) transformation. The latter is produced using a one level filter pass and a symlet of order 5 (Orlov et al. 2008). It also considers an Edge (E) transformation, a standard transformation in digital image processing, which is carried out by approximating the image gradient with the Prewitt operator (Prewitt 1970). Finally, also transformations of transformations are analysed including the Chebyshev transformation of the image's Fourier transform, C(F), as well as W(F), F(W), F(C), C(W), F(E) and W(E).

A large set of features, in total 2919, is extracted from the images and its transforms. We list them in table 1, where we also show from which of the numerous image transformations a feature is derived from. The total feature vector can be divided into feature families, each of which with its own set of feature classes. We will describe this hierarchy in the following and strive on the description of each individual feature. For more information on specific image features, we refer the interested reader to the literature provided by the references in the following.

The first family of features is pixel statistics. The simplest class of such features that `WND-CHARM` derives are simple pixel intensity statistics and consists of mean, median, standard deviation, minimum and maximum of the image's intensity values. This provides five features per image and image transform. Slightly more complex are multi-scale histograms of the intensity distribution using three, five, seven or nine bins, respectively. The values for each bin add 24 features per image and transform. The first class of features which contains information on the spatial correlation of intensity values are the combined moments. Mean, standard deviation, skewness and kurtosis are calculated from all intensity values that fall within a stripe along the x-axis through the image centre with a width that is half the total image height. This stripe is then rotated by 45, 90, and 135 degrees around the centre, creating another three sets of first moments and sampled into three-bin histograms. A total of 48 features in this class is the result. More abstract is the Gini coefficient of the image, which was originally defined for economical studies, but which is now often used in astrophysical applications (e.g. Florian et al. 2016). The full definition can be found in (Abraham et al. 2003), in short, the Gini coefficient describes with a single number the level of discrepancy of an image's intensity value distribution from a perfectly equal intensity distribution.

The second family of feature classes derives from polynomial decompositions of the image intensities and some of its transforms. The first class of these decompositions are based on the Chebyshev coefficients from an order $N = 20$ Chebyshev transform. The values of the coefficients are used to fill a 32-bin histogram, providing the 32 image features of this class. Another set of 32 features comes from the coefficient histogram of an order $N = 23$ Chebyshev-Fourier transform (Orlov et al. 2006) of the image. After a Zernike decomposition of the 2D image (Teague 1980), the 72 first Zernike coefficients create another class of features. Finally, the radon transformations (Radon 1917) along lines with inclination angle 0, 45, 90 and 135 degrees are binned into 3-bin histograms which provide another twelve features per image and image transformation.

Image textures, the third feature family, describe the spatial correlation of intensity values. `WND-CHARM` applies Gabor filters (e.g. Fogel & Sagi 1989) for seven frequencies and using a Gaussian harmonic function as a kernel following (Grigorescu et al. 2002). This provides the seven Gabor texture features which are defined as the area, occupied by the Gabor-transformed images (Orlov et al. 2008). Tamura textures describe contrast, coarseness and directionality of an image. We refer the interested reader to Tamura et al. (1978) for a full definition of these properties. Contrast and directionality are two of the Tamura features used in this analysis, coarseness sum and a 3-bin coarseness histogram complete the set of six features in this class. Another, quite specific, feature class are Haralick textures, which are described in Haralick et al. (1973). As many texture features, they are calculated from the spatial grey-level dependence matrix (SGLDM) of the image which encodes how many times per image specific intensity value pairs are located at a certain distance and along a certain direction. Haralick textures are the fundamental statistical properties of this matrix and contain important features used in digital image processing such as the angular second moment (ASM), the image contrast, correlation and entropy. The 28 properties of the SGLDM defined in the appendix of Haralick et al. (1973) are the Haralick features used by `WND-CHARM` and in this study. In addition to image textures based on the SGLDM, Wu et al. (1992) proposed to characterise the roughness of an image by using a fractional Brownian motion model on multiple scales. The 20 image features coming from the fractal model class implemented in `WND-CHARM` complete the texture family of features

Probably more related to the human perception of images is the last family of features which works on the raw pixel data only and tries to detect distinct objects in the image. The edge feature class applies a Prewitt operator (Prewitt 1970) to approximate the image gradient. The mean, median, variance and 8-bin histogram of both magnitude and directionality of the gradient provide 22 image features. They are complemented by the total number of edge pixels, their direction homogeneity and a 4-bin histogram of all possible differences between the directionality histogram bins, that were mentioned earlier. The final class of object features are Otsu features. They are calculated after applying a global Otsu threshold (Otsu 1979) to the image, thereby converting the image to binary. Basic statistics of all eight-connected (connected to edge or corner) objects in this mask are then collected including, their abundance, Euler number and centroid (both coordinates). In addition, minimum, maximum, mean, median variance and a 10-bin histogram are calculated for area and distance to the centroid of all objects. In total 34 features. The same procedure is applied to the inverse Otsu binary image, adding a final set of 34 features.

We refer the reader to table 1 for an overview of all features and to which group of image transforms they are applied to.

Table 1. Image features used in this analysis.

Family	Class	Features	Input	Reference
Pixel statistics	Combined moments	48	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	–
	Gini coefficient	1	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	Abraham et al. (2003)
	Multiscale histograms	24	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	–
	Pixel intensity statistics	5	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	–
Polynomial decomposition	Chebyshev coefficients	32	raw, F, W, C, F(W), E, F(E), W(E)	–
	Chebyshev-Fourier coefficients	32	raw, F, W, C, F(W), E, F(E), W(E)	Orlov et al. (2006)
	Radon coefficients	12	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	Radon (1917)
	Zernike coefficients	72	raw, F, W, C, F(W), E, F(E), W(E)	Teague (1980))
Textures	Fractal analysis	20	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	Wu et al. (1992)
	Gabor	7	raw	Fogel & Sagi (1989)
	Haralick	28	raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	Haralick et al. (1973)
	Tamura	6	raw raw, F, W, C, C(F), W(F) F(W), F(C), C(W), E, F(E), W(E)	Tamura et al. (1978)
Objects	Edge features	28	raw	Prewitt (1970)
	Otsu object features	34	raw	Otsu (1979)
	Inverse Otsu object features	34	raw	Otsu (1979)

2.5 Classification

The computer vision characterisation algorithm produces a feature vector F for each image, from both training and test sets. The number of features in this vector shall be $|F| = M$. One can now derive a simple classification scheme based on a metric for the distance between a test image to the training set classes in the M-dimensional feature space or a subset of it. The total number of classes in the training set shall be N. Not every feature is equally informative in discriminating one training set class from another. This is why we assign weights w_f to each single feature f following a Fisher discriminant criterion (e.g. Bishop 2006)

$$w_f = \frac{\sum_{c=1}^{N} \left(\langle F_f \rangle - \langle F_f^c \rangle \right)^2}{\sum_{c=1}^{N} (\sigma_f^c)^2} \frac{N}{N-1}, \tag{1}$$

where $\langle F_f \rangle$ is the mean of all values of the feature f among all images in the training set and $\langle F_f^c \rangle$ is the mean of all values of a specific feature f within a single class c of the training set and $(\sigma_f^c)^2$ is its variance within that class. In this, all variances are calculated after the feature values are normalised to fall into the interval $[0, 1]$.

Once all feature weights are calculated from the training set, we can calculate the feature space distance of a single test image feature vector F from a training set class c. One example of such a distance metric which is readily implemented in `WND-CHARM` is the weighted nearest neighbour (WNN) distance

$$d_{\mathrm{WNN}}^c = \min_{T \in T_c} \sum_{f=1}^{M} w_f \left(F_f - T_f \right)^2, \tag{2}$$

where T is a feature vector from the training set, T^c is the set of all feature vectors in the training set that belong to class c and w_f are the feature weights as defined by equation 1. Another feature distance metric, which is also readily implemented in `WND-CHARM`, is the weighted neighbour distance (Orlov et al. 2008)

$$d_{\mathrm{WND}}^c = \frac{\sum_{T \in T_c} \left[\sum_{f=1}^{M} w_f \left(F_f - T_f \right)^2 \right]^p}{|T^c|} \tag{3}$$

where $|T^c|$ is the number of training images in class c and p is a free parameter. The main difference between the WNN and WND distance metric is that WNN defines the distance between a feature vector F and a class T^c by taking into account only the smallest of all distances (or generally k-smallest distances) between F and all $T \in T^c$. In contrast the WND metric takes into account all distances to the images of the training set of class c, but suppresses the larger contributions by means of the exponent p. It was shown in Orlov et al. (2008) that the classification accuracy is not particularly sensitive to p as long as it is smaller than -4. In the `WND-CHARM` package it is fixed to -5. Orlov et al. (2008) also showed that the WND metric gives slightly better results than WNN in biomedical image classification problems.

Once we decided on a distance metric we can define the similarity of a given test image described by its feature vector F to the classes c defining the training set as

$$S_F^c = \left(d_F^c \sum_{i=1}^{N} (d_F^i)^{-1} \right)^{-1} \tag{4}$$

This conveniently assigns for each test image–class pair a similarity $S_F^c \in [0, 1]$. The classification is then performed by assigning the class with the highest similarity to the test image.

At this point it has to be mentioned that other classification schemes are possible, e.g. based on machine learning (including support vector machines, neural networks or Gaussian processes) and potentially implemented as a Bayesian scheme (e.g. Bishop 2006). We chose a simple feature distance based scheme since we wanted to retain a maximum number of usable features while achieving acceptable runtimes in this exploratory study. Other classification schemes shall be explored in future work, the general methodology is not limited to our rather simple current approach.

A more detailed description of `WND-CHARM`, the software used for image characterisation and classification in this work is given in Orlov et al. (2008). A specific description of the software package and its usage is given by Shamir et al. (2008) and examples for its applications in bio medical and astronomical classification problems can be found in Shamir et al. (2010) and Schutter & Shamir (2015).

3 RESULTS

The methodology described in sections 2.4 and 2.5 is applied to the simulated data described in section 2.1. Each of the eight classes in the training set contains 4000 halo images and comes with a set of 200 halo images for testing the classifier. The test sets are available for six different signal to noise levels. Every class has six different smoothing lengths applied to it and is analysed at three different redshifts.

In this section we first identify the most severe degeneracies between the models and reduce the number of classes to four. We then report on the classification success rates as a function of external parameters such as signal to noise in the images and applied smoothing scale and as a function of physical parameters such as redshift or halo mass.

3.1 Model degeneracies

Some of the modified gravity models are expected to give very similar results given the choice of their parameters. We hence identify the most severe degeneracies in order to discard models from this initial study which are almost impossible to distinguish from one another based on image morphology characterisation only. As a first step we therefore run a series of classifications with only two model classes and cycle through all possible combinations of model pairs. The success rates of these classification runs, so the ratio between the correctly classified haloes in a certain class and the total number of tested haloes is reported in table 2. In this first set of classifications we choose the noise-free image samples and apply no smoothing. To perform the classification we use the WND distance metric on the full feature vector containing 2919 image features. This set up is fiducial for classification runs and used if no further comments are made.

In this case, the expected success rate for a classification by random pick is 0.5 and the table highlights in bold font the rates for all model combinations which are smaller then 0.75. As expected, the *f6* and *symmA* models differ only mildly from *lcdm*. The degeneracy between *symmA* and *f6* is even stronger. Interestingly, the three DGP models are particularly distinct from e.g. *lcdm*, but they are quite degenerate among themselves. As a main result of this analysis we drop in the following the *f6*, *symmA*, *dgp05* and *dgp12* models. From now on we will only focus on a single model per modified gravity class namely *lcdm*, *f5*, *symmB* and *dgp56*. The choice of *lcdm* is clear since it is the baseline cosmology. *f5* and *symmB* are chosen since they eliminate the worst degeneracies to *lcdm* within their gravity model class. The choice of *dgp56* is more arbitrary but is based on the fact that all DGP models seem similarly degenerate to the other gravity models and *dgp56* should give the smallest differences compared to *lcdm*. We hence picked it as a worst case scenario. We will, however, revisit this choice later in the analysis. For completeness, we show a number of key analyses for the full set of eight models in appendix A.

3.2 General classification success rates

The main classification result for the remaining four models is shown in table 3. We report there the fiducial classification run at $z = 0$, with no added noise and with no smoothing applied. Furthermore, the full feature vector with 2919 entries was used to calculate the similarities. Each main row of the table represents a class of test sets. The first line in each of these rows shows the number of images which were classified as a member of the respective training set class indicated by the column. The second line in each model row shows the average similarity value (see equation 4) for all test images of the current test class to every training set class. The last line in the model row is just this similarity value normalised to the similarity of the correct test image class. Finally, also the overall success rate per model is shown in the last column of each model row and at the very bottom of the table we report the total success rate of the classification run.

The overall success rate is 61%, far better than the expected result of 25% for a classification by random pick. The results for the individual models are quite different. 122 out of 200 *lcdm* haloes were classified correctly and by looking at the average similarities, the model appears quite distinct from *dgp56*; it is mildly degenerate with *f5* with 26 mis-classifications and more severely degenerate with *symmB* with 42 mis-classifications out of 200 haloes. *f5* appears to be the most difficult model to classify. Although its success rate of 46% still clearly favours the correct answer, there is a strong degeneracy with *symmB* and a mild degeneracy with *lcdm*. Again, there are almost no mis-classifications of *f5* images as *dgp56*. The best result is indeed seen for the *symmB* classification, with 142 our of 200 haloes or 71%. A small and equal trend to mis-identify *symmB* haloes as *lcdm* or *f5*, respectively is apparent. Also *dgp56* is classified solidly with a 65% success rate. Most interestingly there is a pronounced degeneracy with *lcdm*, something that is not seen vice-versa.

3.3 Classification success rate dependence on external parameters

In contrast to the fiducial classification, figure 1 shows the overall classification success rate while varying some external parameters. These parameters include the level of white noise that was added to the test images, the level of smoothing that was applied to the training and test images and the

Table 2. Classification success rates for all possible direct model pair comparisons.

	lcdm	f5	f6	symmA	symmB	dgp05	dgp12	dgp56
lcdm	*	0.83	**0.70**	**0.66**	0.84	0.79	0.78	0.80
f5	*	*	0.8	0.76	0.78	0.83	0.82	0.82
f6	*	*	*	**0.53**	0.77	0.79	0.76	0.79
symmA	*	*	*	*	0.79	0.79	0.77	0.80
symmB	*	*	*	*	*	0.87	0.82	0.84
dgp05	*	*	*	*	*	*	**0.61**	**0.72**
dgp12	*	*	*	*	*	*	*	**0.62**

Table 3. Classification matrix for four models using the full feature vector at no added noise or smoothing.

	lcdm	f5	symmB	dgp56	success rate
	122	26	42	10	
lcdm	0.3698	0.2535	0.2422	0.1345	0.61
	1.00	0.69	0.66	0.36	
	34	92	66	8	
f5	0.2467	0.3579	0.2904	0.1049	0.46
	0.69	1.00	0.81	0.29	
	27	25	142	6	
symmB	0.2374	0.2801	0.3829	0.0996	0.71
	0.62	0.73	1.00	0.26	
	40	12	19	129	
dgp56	0.1709	0.1414	0.1361	0.5516	0.65
	0.31	0.26	0.25	1.00	
all					0.61

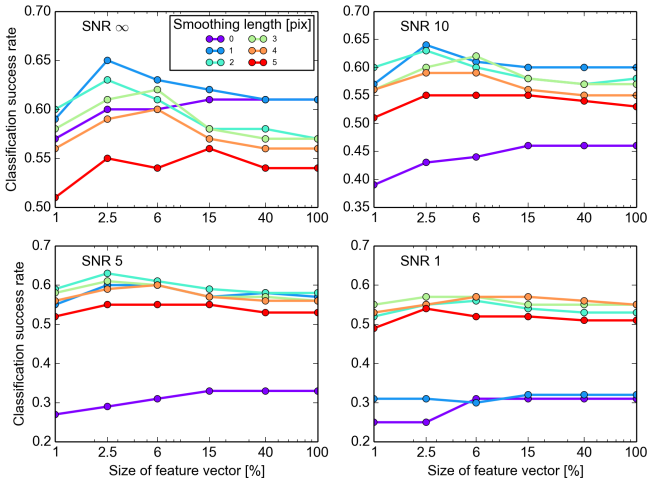

Figure 1. The overall classification success rate as a function of several external parameters. The four panels show classification runs with different levels of noise applied to the test images, as indicated in the top left corner of each panel. The colour-coded lines show different levels of smoothing applied to both training and test images and the x-axis in each panel refers to the percentage of fisher-weight ranked image features that were used for the classification, where 100% is 2919 features.

number of parameters that were used to perform the classification. Figure 1 shows a number of expected results. Firstly, the classification success rate drops quickly in the presence of increasing noise levels. From the familiar ∼60% in the absence of noise, down to less than 30% for very noisy images with a signal to noise ratio of 1. However, this trend can readily be fixed with the application of mild levels of smoothing. A rather small smoothing scale of one pixel is enough to bring the classification rate back to about 60%. In the extreme case of a signal to noise ratio of 1, smoothing on a scale of three pixels is needed to achieve this. While smoothing is crucial in the presence of noise, it only mildly affects the classification rate in the absence of noise. Even for a rather extreme smoothing scale of five pixels, the success rate drops by only 7%. Most interestingly, the best overall classification result is not achieved for the fiducial case, but for a smoothing scale of one pixel, while using only the top 2.5% of the elements in the fisher-weight ranked feature vector, so the 73 most discriminating features.

Figure 2 is similar to figure 1 but shows the overall classification success rate for both, the WND and WNN distance metric. We confirm the trend reported by Orlov et al. (2008), that the WND is generally superior to WNN. The only exception are cases of low signal to noise in the images and in the absence of smoothing (bottom-left panel of figure 2). However, this configuration delivers generally bad results as discussed earlier and highlighted in figure 1. It seems that in this case the WND strategy of considering the distances to all training images of a class can indeed introduce significant biases in the presence of noise. We will investigate this behaviour in more detail in the discussion of the following section.

3.4 Classification success rate dependence on physical parameters

In order to study the classification success rate as a function of physical properties we further divide each of the four model classes into four mass bins of equal object number. The first bin contains the 50 most massive systems in the sample, the second mass bin the $50^{\rm th}$ to $100^{\rm th}$ most massive system and so on. The overall classification results and the results for each individual model are reported in table 4. As can be seen there, the classification success can indeed increase by up to 10% when compared to the fiducial single mass bin run. The only exception is the first, most massive mass bin, where we have a wide mixture of masses due to the steep drop of the mass function at the high-mass end. For the other mass bins, where the mass separation is more effective, the increase is more pronounced.

Table 5 shows the overall and individual model classification success rate in the three available redshift bins. We

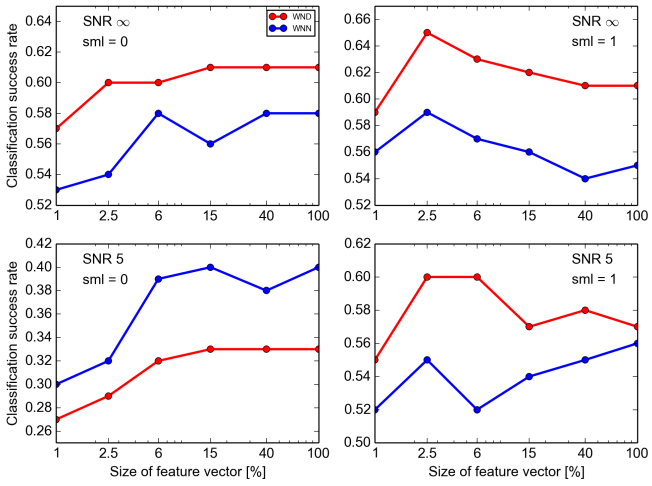

Figure 2. The overall classification success rate as a function of distance metric. Similar to figure 1, the four panels show different levels of noise in the test images, but in this plot they also encode a fixed smoothing scale that was applied to both, test and training set images. The size of the feature vector is given as percentage relative to the full 2919 features by the x-axis. The WND distance metric (equation 3) is shown in red, the WNN distance metric (equation 2) is shown in blue.

Table 4. Classification success rate in bins of halo mass

mass bin	1	2	3	4	all
lcdm	0.64	0.74	0.64	0.72	0.61
f5	0.52	0.60	0.50	0.54	0.46
symmB	0.66	0.74	0.74	0.66	0.71
dgp56	0.68	0.74	0.90	0.80	0.65
all	0.63	0.71	0.70	0.68	0.61

see a continuous decrease with increasing redshift, related to the fact that at high redshift the distinguishing effects of each model had less time to act on the morphology. In the course of this redshift dependence study and as an additional exercise we attempted a classification by redshift. We take all images of the *lcdm* model and group them into redshift classes. We then use these classes and the familiar characterisation and classification scheme to determine the redshift for each of the respective test images. The results are reported in table 6, which follows the same layout as table 3. The success rate at low redshift is almost 100% and we find a good 80% for our highest redshift class. The intermediate redshift class gives the worst result since it is bracketed by two instead of only one other model and is hence more prone to degeneracies. Still, a 74% success rate shows the ability to sort objects by redshift based on their image morphology only.

4 DISCUSSION

We now revisit and elaborate on the more interesting results of section 3.

Table 5. Classification success rate at different redshifts

z	0.0	0.5	1.0
lcdm	0.61	0.51	0.46
f5	0.46	0.58	0.52
symmB	0.71	0.55	0.39
dgp56	0.65	0.64	0.55
total	0.61	0.57	0.48

Table 6. Classification matrix for redshift classes.

	z00	z05	z10	success rate
z00	193 0.8638 1.00	7 0.1268 0.15	0 0.0094 0.01	0.97
z05	26 0.1558 0.25	148 0.6217 0.35	26 0.2225 1.00	0.74
z10	1 0.0120 0.02	40 0.2554 0.35	159 0.7326 1.00	0.80
all				0.83

Table 7. Classification matrices in the absence of noise with optimal choice of smoothing scale (one pixel) and feature vector size (∼70).

	lcdm	f5	symmB	dgp56	success rate
lcdm	129 0.4704 1.00	27 0.2065 0.44	31 0.1935 0.41	13 0.1295 0.28	0.65
f5	39 0.2167 0.51	104 0.4287 1.00	53 0.2672 0.62	4 0.0874 0.20	0.52
symmB	32 0.1701 0.34	22 0.2497 0.50	140 0.5028 1.00	6 0.0773 0.15	0.70
dgp56	26 0.1492 0.23	9 0.1010 0.15	17 0.0979 0.15	148 0.6519 1.00	0.74
all					0.65

4.1 Optimal parameter choices

As figure 1 implies, there is a benefit of applying a mild one pixel smoothing to images, even if they are noise free, and then limit the number of features to be used for the classification to about 2% of the total feature vector. With the optimal choice of smoothing and feature vector length, we obtain the classification result that is shown in table 7. The overall classification success rate raises slightly by 4% to 65%. The largest individual model increases in the success rate are seen for *dgp56*, from 65% to 74% and for *f5*, from 46% to 52%. It seems that the combination of smoothing and feature reduction removes a number of degeneracies to other models, mostly to *symmB* for *f5* and to *lcdm* for *dgp56*.

Table 8. Classification success rates for different settings of external parameters and depending on the distance metric used for classification.

	WND	WNN
SNR = 5, no smoothing		
lcdm	0.18	**0.41**
f5	0.0	**0.15**
symmB	0.27	0.27
dgp56	**0.88**	0.77
total	0.33	**0.40**
SNR = 5, 1 pixel smoothing		
lcdm	**0.61**	0.50
f5	0.40	0.40
symmB	**0.70**	0.59
dgp56	**0.71**	0.70
total	**0.60**	0.54

4.2 Response to noise

The WND distance metric turned out to be the best choice in most cases. The only exception that shows up in figure 2 is the low signal to noise case without subsequent smoothing. To study this case in more detail we show the success rate for individual classes in table 8. Most interestingly, the WND metric fails completely to classify successfully the *f5* model in the absence of smoothing, but excels over WNN in the case of *dgp56*. It appears that the strategy of WND to take into account all training images of a class, even while suppressing the contribution of the more distant samples in feature space, introduces a severe bias in the treatment of *f5* images. When looking into the results in more detail, it shows that 128 out of the 200 *f5* images get classified as *dgp56* in this case. This is a first hint of irregularities with the *dgp56* model in the presence of noise and in the absence of smoothing, which we shall investigate further.

In order to do so, we show in figure 3 the success rate per model as a function of signal to noise ratio (SNR) and with different levels of smoothing applied. The bottom two panels of the figure, with strong smoothing, show the typical success rate hierarchy of models, which is largely constant over the full SNR range. With no or weak smoothing applied, however, a pattern emerges. The success rate for *dgp56* peaks sharply with an ideal success rate at SNR = 2, which slightly drops for the lowest SNR images. The success rate of all other models vanishes at this *dgp56* SNR sweet spot. This trend is clarified by figure 4, which, as a function of the same parameters shows the relative mis-classification rate per model. The combination of the two figures tells us that at SNR = 2 and without smoothing, every image gets classified as *dgp56*. This trends shifts to SNR = 1 if we apply a one pixel smoothing. This suggests that image features of the *dgp56* training set can be mimicked by noise in the test images. This is good and bad. It is good because it already hints towards that *dgp56* images seem to contain a lot of small-scale structure, but it is also very alarming since it clearly shows that our analysis can be largely biased towards specific models in the presence of noise. Luckily, the described effects can be easily alleviated by the application

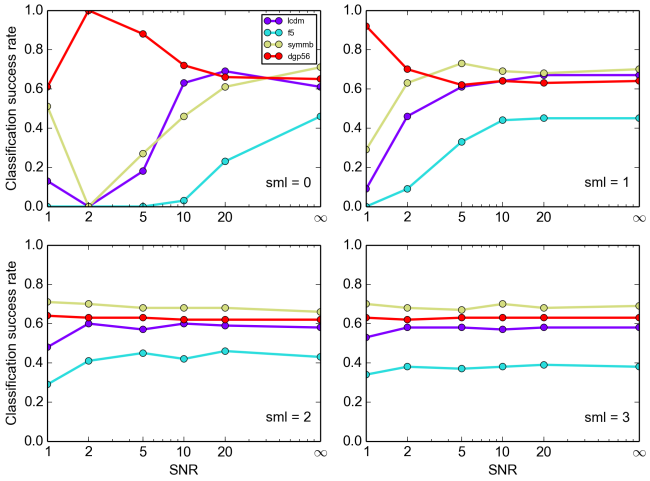

Figure 3. Classification success rate for individual models as a function of signal to noise of the test images. The four different panels refer to four different smoothing scales applied to training and test data, indicated by the label in the bottom-right corner of each panel.

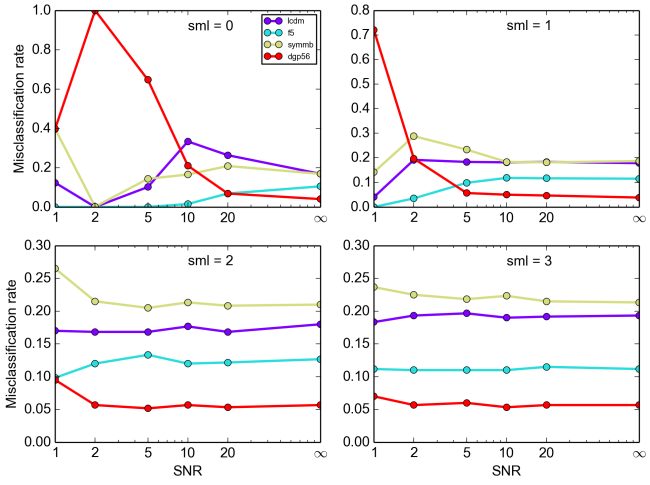

Figure 4. The mis-identification rate for all models as a function of the same parameters as in figure 3.

of smoothing. However, our noise model was particularly simple and the addition of more realistic noise might give way to more subtle effects.

4.3 Dependence on σ_8

All of our simulations have exactly the same initial conditions. One consequence of this is that the amplitude of matter fluctuations today, σ_8, will be different: for *lcdm* we have $\sigma_8 = 0.8$, for *dgp12*, *f5* and *symmB* we have $\sigma_8 = 0.85$ while for *dgp56*, *f6* and *symmA* we have $\sigma_8 = 0.82$. *dgp05* produces $\sigma_8 = 0.89$.

The difference in structure formation between modified gravity models and ΛCDM will be degenerate with a change

Table 9. Direct model on model comparison for two *lcdm* models, with different σ_8 parameters.

	$\sigma_8 = 0.8$	$\sigma_8 = 0.85$	success rate
$\sigma_8 = 0.8$	192 0.7829 1.0	8 0.2171 0.28	0.96
$\sigma_8 = 0.85$	56 0.2515 0.34	144 0.7485 1.00	0.72
all			0.84

Table 10. The same comparison as table 9 but for SNR = 5 and a smoothing scale of 2 pixels.

	$\sigma_8 = 0.8$	$\sigma_8 = 0.85$	success rate
$\sigma_8 = 0.8$	187 0.7722 1.0	13 0.2278 0.29	0.94
$\sigma_8 = 0.85$	60 0.3016 0.43	140 0.6984 1.00	0.70
all			0.82

of this amplitude, i.e. we can mimic some of the effects of a modified gravity model at a given redshift by simply running a ΛCDM simulation with a higher σ_8. This is especially true for the DGP models for which the modifications of gravity, in the absence of screening, is equivalent to a time-dependent Newtonian constant $G_{\rm eff}(a) = G(1 + C(a))$. The presence of a screening mechanism, which leads to general relativity being recovered inside the most massive haloes, will act to reduce this degeneracy.

We investigated the response of the classification results to different values of σ_8 by running a new *lcdm* simulation with $\sigma_8 = 0.85$ instead of $\sigma_8 = 0.8$. We then directly compare the two *lcdm* models and show the classification matrices in table 9. The classification is indeed quite sensitive to this change in parameters. The result is confirmed in table 10, for a lower SNR = 5 and with a two pixel smoothing applied. On the one hand, this is very encouraging since it shows that the method is sensitive to changes in cosmological parameters for the same model of gravity. Hence, the computer vision approach can be used to constrain cosmological parameters from observations when using numerical simulations on a parameter grid (e.g. Heitmann et al. 2016, for a recent suitable simulation suite) as training sets. But on the other hand, it also shows us that some some effects we were seeing earlier indeed stem from a different σ_8 seen in the different models. We therefore run another set of classifications, where we exchange the *dgp56* model with the *dgp12* model, which produces a different σ_8. We show the results in table 11 and figure 5. There we see that the overall classification success rate is slightly worse, due to the fact that *dgp12* is slightly more degenerate with other models. However, all major trends that we have seen earlier are effectively unchanged, which suggests that the change in σ_8 between the simulations is indeed not the main driver for the ability to distinguish the models. A similar classification as the one underlying table 11 is shown in appendix A, where the baseline *lcdm* model is switched in place with the one the produces the higher σ_8.

4.4 Image features

To look more into the specific images features of each model, we take a closer look at the similarity parameter defined by equation 4 for individual test images. In the first four rows of figure 6 we show those five test images per class that scored the highest similarity values to their true model, respectively. The bottom row shows the most ambiguous images, where the distribution of similarities to all classes is most uniform[3]. The images in the four top rows can hence we interpreted as prototype images, best representing the features of their respective training set class. For the *lcdm* model we see mostly isolated haloes with a quite elliptical core, surrounded by a number of small satellites. The fourth *lcdm* image in figure 6 shows a merger, where each merging component in the field is again a elliptical core halo with a number of satellites. At least by visual inspection, the typical *f5* halo has a rounder core and is surrounded by more massive substructure. The same is true for *symmB* but the overall structure of the images looks more complicated, with either very elliptical or extremely complex cores which can show several peaks. As expected from the earlier discussion, the *dgp56* images look most distinct with very pronounced filamentary structure around the halo and many coarse, low-intensity features along these filaments. We assume that it is this filamentary structure which can be easily mimicked by noise and bias the classification towards this model. The most ambiguous haloes are usually some special situations where we have several very massive and distinct structures on a single field or particularly complicated merger situations.

Table 11. Classification matrices with the *dgp12* model.

	lcdm	f5	symmB	dgp12	success rate
lcdm	123 0.3537 1.00	27 0.2556 0.72	42 0.2530 0.72	8 0.1377 0.39	0.62
f5	36 0.2555 0.75	90 0.3424 1.00	68 0.2900 0.85	6 0.1121 0.33	0.45
symmB	29 0.2514 0.70	24 0.2767 0.77	141 0.3596 1.00	6 0.1124 0.31	0.71
dgp12	46 0.1799 0.35	13 0.1498 0.29	25 0.1544 0.30	116 0.5160 1.00	0.58
all					0.59

Obviously, the descriptions above are very subjective and may miss subtle features. In fact our aim to use computer vision based classifiers is motivated in avoiding such

[3] The five images were chosen such to have the minimum variance in their four similarity values.

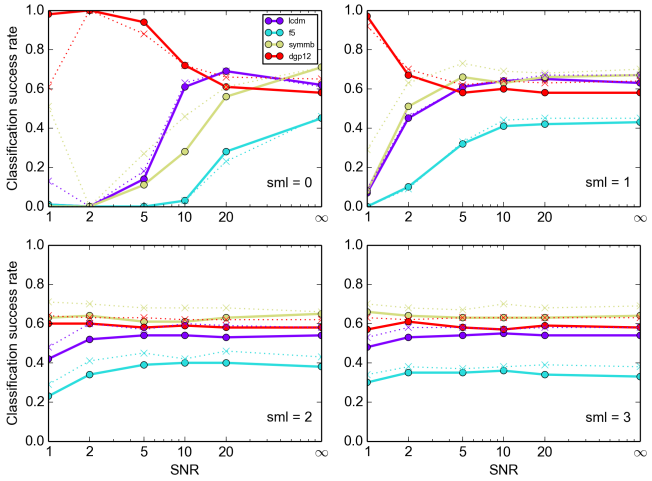

Figure 5. This figure is the same as figure 3, but with the *dgp56* model switched in place with a *dgp12* model. For comparison, the former results of figure 3 are shown by the thinner, dotted lines.

descriptions but to quantitatively characterise images. The interpretation of 2D halo morphology for the modified gravity models of table using computer vision is listed in table 12, which shows the 25 image features with the highest fisher discriminants for the classification run summarised in table 7. This particular list is derived from training sets smoothed on a one pixel scale, which produced the best results, but we have verified that there are only marginal changes in feature fisher scores between different smoothing scales. For completeness we list the remaining 48 features that were used in this analysis in appendix B. Some of the top-ranked features are easy to interpret in terms of halo morphology. For example, the first Haralick texture is the angular second moment, which is a measure of the complexity of an image. The fact that this feature is most descriptive on the Edge transform of the image is not surprising since this transformation just filters out the actual structure in the image and hence reduces the noise. However, most of the other main features are more difficult to interpret directly, such as the Zernike-coefficients of the Fourier transform, the Zernike-coefficients of the Fourier transform of the Wavelet transform and Tamura textures of the Chebyshev transform. Another very discriminating feature seems to be the Gini coefficient for different transformations, another measure that describes the homogeneity of an image. To get a bigger picture we list in table 13 the 15 most significant feature classes ordered by the mean fisher score within the class. Also minimum, maximum and standard deviation of the fisher scores within the family are given in the table. Again, we see that the Gini coefficient for different transformations, together with Haralick and Tamura textures and Zernike coefficients are most relevant. One should not over-interpret the actual ranking within that list though since the number of features in each family is varying strongly. For example, the single features with the highest Fisher score belong to the families of Haralick textures and Zernike coefficients, but they are hidden within the mean of the feature family in this table. While it would be excessive to list the fisher scores for all

Table 12. The most discriminating image features according to their fisher discriminant.

Rank	Name	Weight
1	Haralick Textures (Edge ()) [0]	0.4908
2	Zernike Coefficients (Fourier ()) [19]	0.4173
3	Zernike Coefficients (Fourier ()) [26]	0.4091
4	Zernike Coefficients (Fourier ()) [41]	0.402
5	Zernike Coefficients (Fourier ()) [54]	0.3961
6	Zernike Coefficients (Fourier ()) [28]	0.3958
7	Tamura Textures (Chebyshev ()) [2]	0.3917
8	Haralick Textures (Edge ()) [22]	0.3881
9	Haralick Textures (Edge ()) [18]	0.3838
10	Zernike Coefficients (Fourier ()) [5]	0.3779
11	Zernike Coefficients (Fourier (Wavelet ())) [19]	0.3723
12	Zernike Coefficients (Fourier ()) [71]	0.372
13	Zernike Coefficients (Fourier (Wavelet ())) [28]	0.3718
14	Gini Coefficient (Fourier (Wavelet ())) [0]	0.3661
15	Zernike Coefficients (Fourier ()) [69]	0.3655
16	Zernike Coefficients (Fourier (Wavelet ())) [69]	0.3648
17	Haralick Textures (Edge ()) [10]	0.3634
18	Zernike Coefficients (Fourier (Wavelet ())) [54]	0.3618
19	Zernike Coefficients (Fourier (Wavelet ())) [5]	0.3612
20	Zernike Coefficients (Fourier (Wavelet ())) [41]	0.357
21	Zernike Coefficients (Fourier (Wavelet ())) [39]	0.3553
22	Gini Coefficient (Fourier ()) [0]	0.3538
23	Haralick Textures (Edge ()) [15]	0.3508
24	Zernike Coefficients (Fourier ()) [67]	0.3482
25	Haralick Textures () [4]	0.3462

2919 features[4], but we list in appendix B the fisher score statistics of the remaining 109 feature families.

5 SUMMARY AND OUTLOOK

This work introduces computer vision characterisation of 2D images of dark matter surface mass density distributions to classify different models of structure formation. We used a set of numerical simulations, all run with identical initial conditions, but with different underlying models of gravity. Our simulations included a standard ΛCDM model, two $f(R)$ models, two Symmetron models and three DGP models. We extracted the 200 most massive haloes in each of the simulation boxes and produced 4000 images of the surface mass density distribution along 20 different lines-of-sight. On these images we ran a computer vision characterisation algorithm, which extracts up to 2919 image features and then used a simple WND and WNN based classification algorithm to classify a set of 200 test images for each class of models. Both, the computer vision characterisation algorithm and the subsequent classifier are implemented in the publicly available software package `WND-CHARM`. To the test images we applied several levels of uncorrelated noise and probed the effects of smoothing on both, the training and test images.

There are degeneracies between certain models when using image morphology for classification. Especially the DGP models are strongly degenerate between themselves,

[4] Full classification data available upon e-mail request from the authors.

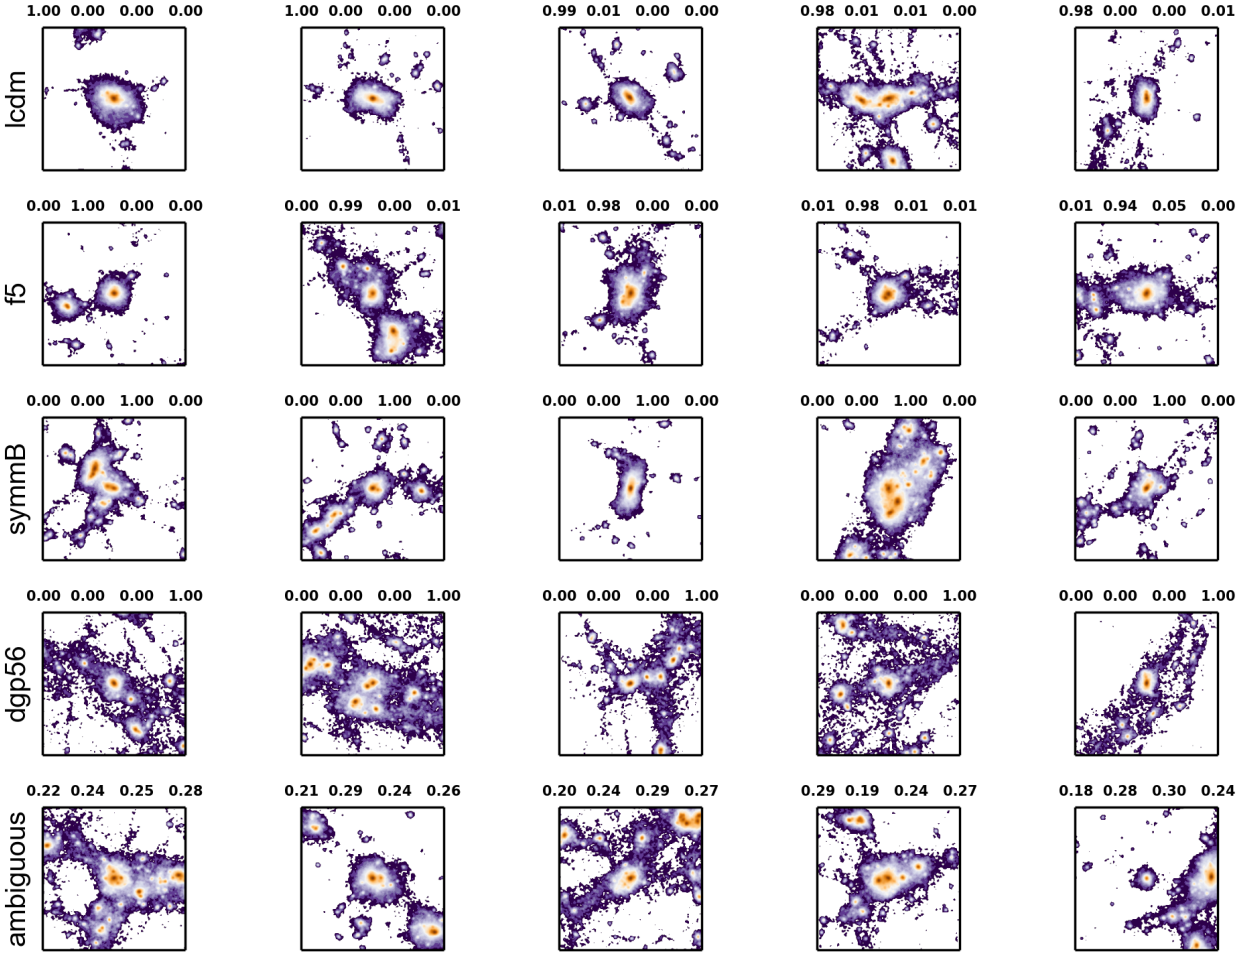

Figure 6. Test images with peculiar similarity values. The first four rows show the five images per model class with the highest similarity to their true model class. The numbers above each thumbnail image show similarity value to the *lcdm*, *f5*, *symmB* and *dgp56* class of the training set, respectively. The bottom row shows the five most ambiguous test images, meaning that they are classified as similarly close to all training set classes.

Table 13. Fisher discriminant statistics of feature classes ordered by the mean.

Rank	Name	Min	Max	Mean	Std. dev.	#
1	Gini Coefficient (Fourier (Wavelet ()))	0.3661	0.3661	0.3661	–	1
2	Gini Coefficient (Fourier ())	0.3538	0.3538	0.3538	–	1
3	Gini Coefficient (Wavelet (Fourier ()))	0.259	0.259	0.259	–	1
4	Gini Coefficient (Chebyshev (Wavelet ()))	0.2485	0.2485	0.2485	–	1
5	Gini Coefficient (Edge ())	0.2361	0.2361	0.2361	–	1
6	Gini Coefficient (Wavelet (Edge ()))	0.2178	0.2178	0.2178	–	1
7	Tamura Textures (Chebyshev ())	0.01879	0.3917	0.2066	0.151	6
8	Haralick Textures ()	0.001923	0.3462	0.153	0.1126	28
9	Haralick Textures (Edge ())	0.002559	0.4908	0.1489	0.154	28
10	Zernike Coefficients (Fourier ())	0.001196	0.4173	0.1403	0.1268	72
11	Zernike Coefficients (Fourier (Wavelet ()))	0.0004068	0.3723	0.138	0.1139	72
12	Gini Coefficient ()	0.1269	0.1269	0.1269	–	1
13	Tamura Textures (Chebyshev (Wavelet ()))	0.0	0.2066	0.1167	0.09447	6
14	Gini Coefficient (Chebyshev ())	0.1046	0.1046	0.1046	–	1
15	Tamura Textures (Fourier ())	0.0	0.2337	0.083	0.1123	6

while Symmetron and $f(R)$ models show only mild deviations from ΛCDM. To quantify these degeneracies we ran model-on-model classifications looping through every possible model pair and report the results in section 3. After this test, we reduced the number of models to four, with only a single model per gravity family. Using the classifier now on these four models shows a typical classification success rate of about 60%. While this is a good success rate, it is not robust enough to accurately classify single haloes. We want to highlight, however, that each class of test images was robustly classified correctly as an ensemble, as shown in table 3. This result even holds when using the full set of eight models as shown in table A1. As expected, the method is quite sensitive to the presence of noise, which in our case was just white pixel noise, applied to the test images. The classification success rates drop to the level of 30% as shown in figure 1, which is only slightly better than a characterisation by random pick. But figure 1 also shows that this sensitivity to noise can easily be remedied by the application of mild smoothing to the training and test images, while only marginally affecting the overall classification success rate compared to the ideal case where no noise is present and no smoothing is applied.

Several strategies can be applied to increase the classification success rate, e.g. by smoothing the training sets and restricting the classification to only a subset of the top fisher score weighted characterisation features. Another possibility is to use additional information on the halo images, for example the a rough estimate of the mass of the object. With such techniques, overall classification success rates can exceed 70%. While the classification accuracy, as expected, decreases for larger redshifts (table 5), our methodology can be used to classify haloes into redshift classes based on their morphology only (table 6). Since all our simulations where run with the same initial conditions, different gravity models may produce different amplitudes of the density fluctuations today, usually expressed in terms of the parameter σ_8. We have extensively tested that our classification is not solely driven by potential variations of σ_8 in the probed models. The approach is however sensitive to these variations and it could be used to constrain it from observations. These results are summarised in section 4.3 and appendix A.

The presence of noise can trigger strong biases and prefer the classification of certain models over others as we show in section 4. Figure 6 shows that the DGP models tend to produce low pixel intensity filamentary structure around the more massive haloes. This morphological feature can easily be mimicked by the addition of noise and is then incorrectly picked up as a feature by the classifier. Although this is easily mended by the introduction of suitable smoothing, it clearly shows that some models are more susceptible to noise than others. These findings also highlight that the morphology of the different models indeed differ significantly, firstly shown by the ability of our classifier to distinguish them and visually highlighted by figure 6, which shows the most typical prototypes of a certain class, as well as the most ambiguous haloes in the test sample.

As the main finding of this work, we present the most discriminating image features to classify the halo images. Their individual and average feature class fisher discriminant scores are reported in tables 12 and 13, as well as in appendix B. Texture features such as the Tamura and Haralick textures (see section 2.4 for a description) are among the most informative features for classification. While for the Haralick textures, the actual image data or the similar Edge transform is the most descriptive starting point, the Tamura textures are more descriptive in the, to the human perception, more abstract Fourier -or Chebyshev domain. The same is true for the other most descriptive features, being the coefficients of the Zernike decomposition of the image either in Fourier space, or based on the Fourier transform of the Wavelet transform. Finally, the Gini coefficient seems to be a valuable feature in terms of classification when applied to any kind of pixel intensity transformation. In conclusion, this study shows that 2D images of the dark matter distribution can indeed be used to identify the signatures of different models of structure formation, in our case the signatures of different models of gravity. This is most promising since techniques such as gravitational lensing mass reconstructions can deliver such images from observations.

This work is only a first attempt of a computer vision based classification of different models of structure formation. We set it up as an initial proof of concept of the feasibility of this approach. Results seem promising, but we have to point out that our simplistic noise model does not capture all of the complications that a real study of lensing surface mass density maps would entail. The effects of the inclusion of correlated noise deriving from large-scale structure along the line-of-sight (among others Becker & Kravtsov 2011; Hoekstra et al. 2011) needs to be assessed. Furthermore, the systematic effects in the reconstruction of a surface mass density map from lensing data must be mimicked. Such frameworks exist (Meneghetti et al. 2010, 2016) and shall be a next step in this study. It also needs to be verified if such surface mass density maps are indeed needed at all. Computer vision classification could potentially work directly in shear space, a direct lensing observable from galaxy surveys (Bartelmann & Schneider 2001) and a quantity that can be derived, although under more difficulty, from numerical simulations (e.g. Barreira et al. 2016). Also our current classifier based on either WND or WNN distance measures is potentially not the ideal solution. More sophisticated techniques based on machine learning could be applied to the most descriptive feature families, potentially after a dimensionality reduction via a principal component analysis. The overall approach can potentially be extended to much wider fields, probing the morphology of the cosmic web itself and not only of its nodes. With on-going and upcoming wide-field or all-sky galaxy surveys, the necessary input data will be available, but further tests with simulations must verify the robustness of the approach. Finally, in this work we focused only on models of modified gravity to be classified. The halo characterisation framework, however, is fully general and can readily be extended to the classification of different models of dark matter, dark energy interactions, baryonic physics or different global cosmological setups, as we already hinted towards in section 4.3.

ACKNOWLEDGEMENTS

Large parts of this work have been carried out during an Astrophysics summer research programme within the Department of Physics at the University of Oxford. We would

like to thank the organisers for this excellent research opportunity for young students, which enabled QL's substantial contribution to this work. We also want to thank the authors of the excellent `WND-CHARM` algorithm for the development and open access to this software package. JM has received funding from the Marie Curie Actions of the European Union's Horizon 2020 Programme under REA grant agreement number 627288 (WEBMAP). HAW is supported by the European Research Council through 646702 (CosTesGrav).

REFERENCES

Abraham R. G., van den Bergh S., Nair P., 2003, ApJ, 588, 218
Anderson L., et al., 2014, MNRAS, 441, 24
Barreira A., Llinares C., Bose S., Li B., 2016, J. Cosmology Astropart. Phys., 5, 001
Bartelmann M., 2010, Classical and Quantum Gravity, 27, 233001
Bartelmann M., Schneider P., 2001, Phys. Rep., 340, 291
Becker M. R., Kravtsov A. V., 2011, ApJ, 740, 25
Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, ApJ, 762, 109
Bengtsson E., Rodenacker K., 2003, Analytical Cellular Pathology, 24, 1
Betoule M., et al., 2014, A&A, 568, A22
Bishop C. M., 2006, Pattern Recognition and Machine Learning (Information Science and Statistics). Springer-Verlag New York, Inc., Secaucus, NJ, USA
Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2012, MNRAS, 422, 1203
Bozek B., Boylan-Kolchin M., Horiuchi S., Garrison-Kimmel S., Abazajian K., Bullock J. S., 2016, MNRAS, 459, 1489
Bradač M., et al., 2009, ApJ, 706, 1201
Caminha G. B., et al., 2016, A&A, 587, A80
Clifton T., Ferreira P. G., Padilla A., Skordis C., 2012, Phys. Rep., 513, 1
Clowe D., Bradač M., Gonzalez A. H., Markevitch M., Randall S. W., Jones C., Zaritsky D., 2006, ApJ, 648, L109
Dawson W. A., et al., 2012, ApJ, 747, L42
Diego J. M., et al., 2015, MNRAS, 446, 683
Diego J. M., Broadhurst T., Wong J., Silk J., Lim J., Zheng W., Lam D., Ford H., 2016, MNRAS, 459, 3447
Dubois Y., et al., 2014, MNRAS, 444, 1453
Dvali G., Gabadadze G., Porrati M., 2000, Physics Letters B, 485, 208
Florian M. K., Gladders M. D., Li N., Sharon K., 2016, ApJ, 816, L23
Fogel I., Sagi D., 1989, Biological Cybernetics, 61, 103
Gonzalez R. C., Woods R. E., 2007, Digital Image Processing, 3 edn. Prentice Hall
Grigorescu S. E., Petkov N., Kruizinga P., 2002, IEEE Transactions on Image Processing, 11, 1160
Haralick R. M., Shanmugam K., Dinstein I., 1973, IEEE Transactions on Systems, Man, and Cybernetics, SMC-3, 610
Harvey D., Massey R., Kitching T., Taylor A., Tittley E., 2015, Science, 347, 1462
Heitmann K., et al., 2016, ApJ, 820, 108
Heller K. A., Ghahramani Z., 2006, in 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). pp 2110–2117, doi:10.1109/CVPR.2006.41
Hildebrandt H., et al., 2017, MNRAS, 465, 1454
Hinterbichler K., Khoury J., 2010, Physical Review Letters, 104, 231301
Hoekstra H., Hartlap J., Hilbert S., van Uitert E., 2011, MNRAS, 412, 2095
Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, MNRAS, 445, 581
Hu W., Sawicki I., 2007, Phys. Rev. D, 76, 064004
Jauzac M., et al., 2015, MNRAS, 446, 4132
Jee M. J., Hughes J. P., Menanteau F., Sifón C., Mandelbaum R., Barrientos L. F., Infante L., Ng K. Y., 2014, ApJ, 785, 20
Johnson T. L., Sharon K., Bayliss M. B., Gladders M. D., Coe D., Ebeling H., 2014, ApJ, 797, 48
Jullo E., Kneib J., 2009, MNRAS, 395, 1319
Kauffmann G., White S. D. M., Guiderdoni B., 1993, MNRAS, 264, 201
Kimm T., Cen R., Devriendt J., Dubois Y., Slyz A., 2015, MNRAS, 451, 2900
L'Huillier B., Winther H. A., Mota D. F., Park C., Kim J., 2017, preprint, (`arXiv:1703.07357`)
Llinares C., Mota D. F., Winther H. A., 2014, A&A, 562, A78
Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, MNRAS, 439, 300
Massey R., et al., 2007, Nature, 445, 286
Massey R., et al., 2015, MNRAS, 449, 3393
McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, MNRAS, 465, 2936
McGaugh S. S., Lelli F., Schombert J. M., 2016, Physical Review Letters, 117, 201101
Mead A. J., Heymans C., Lombriser L., Peacock J. A., Steele O. I., Winther H. A., 2016, MNRAS, 459, 1468
Medezinski E., Umetsu K., Okabe N., Nonino M., Molnar S., Massey R., Dupke R., Merten J., 2016, ApJ, 817, 24
Melchior P., et al., 2015, MNRAS, 449, 2219
Meneghetti M., Rasia E., Merten J., Bellagamba F., Ettori S., Mazzotta P., Dolag K., Marri S., 2010, A&A, 514, A93+
Meneghetti M., et al., 2016, preprint, (`arXiv:1606.04548`)
Merten J., 2016, MNRAS, 461, 2328
Merten J., Cacciato M., Meneghetti M., Mignone C., Bartelmann M., 2009, A&A, 500, 681
Merten J., et al., 2011, MNRAS, 417, 333
Merten J., et al., 2015, ApJ, 806, 4
Nelson D., Genel S., Pillepich A., Vogelsberger M., Springel V., Hernquist L., 2016, MNRAS, 460, 2881
Newman A. B., Treu T., Ellis R. S., Sand D. J., Nipoti C., Richard J., Jullo E., 2013, ApJ, 765, 24
Oman K. A., et al., 2015, MNRAS, 452, 3650
Orlov N., Johnston J., Macura T., Wolkow C., Goldberg I., 2006, in 3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006.. pp 1152–1155, doi:10.1109/ISBI.2006.1625127
Orlov N., Shamir L., Macura T., Johnston J., Eckley D. M., Goldberg I. G., 2008, Pattern recognition letters, 29, 1684
Otsu N., 1979, IEEE Transactions on Systems, Man, and Cybernetics, 9, 62
Peter A. H. G., Rocha M., Bullock J. S., Kaplinghat M., 2013, MNRAS, p. 600
Planck Collaboration et al., 2016a, A&A, 594, A13
Planck Collaboration et al., 2016b, A&A, 594, A24
Prewitt J., 1970, Picture Processing and Psychopictorics. New York: Academic Press
Radon J., 1917, in Berichte über die Verhandlungen der Königlich-Sächsischen Akademie der Wissenschaften zu Leipzig. Teubner, pp 262–277, doi:doi:10.1109/TMI.1986.4307775, `https://dx.doi.org/10.1109%2FTMI.1986.4307775`
Read J. I., Agertz O., Collins M. L. M., 2016, MNRAS, 459, 2573
Rocha M., Peter A. H. G., Bullock J. S., Kaplinghat M., Garrison-Kimmel S., Oñorbe J., Moustakas L. A., 2013, MNRAS, p. 560
Schaye J., et al., 2015, MNRAS, 446, 521
Schutter A., Shamir L., 2015, Astronomy and Computing, 12, 60
Schwinn J., Jauzac M., Baugh C. M., Bartelmann M., Eckert

D., Harvey D., Natarajan P., Massey R., 2017, MNRAS, 467, 2913
Shamir L., Orlov N., Eckley D. M., Macura T., Johnston J., Goldberg I., 2008, Source Code for Biology and Medicine, 3, 13
Shamir L., Delaney J., Orlov N., Eckley D., Goldberg I., 2010, PLoS Comput Biol 6(11): e1000974
Tamura H., Mori S., Yamawaki T., 1978, IEEE Transactions on Systems, Man, and Cybernetics, 8, 460
Teague M. R., 1980, J. Opt. Soc. Am., 70, 920
Teyssier R., 2002, A&A, 385, 337
Treu T., et al., 2016, ApJ, 817, 60
Umetsu K., et al., 2014, ApJ, 795, 163
Vogelsberger M., Zavala J., Simpson C., Jenkins A., 2014, MNRAS, 444, 3684
Walker M. G., Peñarrubia J., 2011, ApJ, 742, 20
Winther H. A., et al., 2015, MNRAS, 454, 4208
Wu C.-M., Chen Y.-C., Hsieh K.-S., 1992, IEEE Transactions on Medical Imaging, 11, 141
Yavlinsky A., Heesch D., Rüger S., 2006, in Image and Video Retrieval. pp 537–540, http://oro.open.ac.uk/29876/
von Braun-Bates F., Winther H. A., Alonso D., Devriendt J., 2017, J. Cosmology Astropart. Phys., 3, 012

APPENDIX A: FULL SET OF MODELS

This appendix shows some classification results if all eight modified gravity models are used in the analysis. Table A1 shows the same classification success matrices as the four model equivalent in table 3. Although the overall classification success rate and the results for each model are now significantly worse, since no degeneracies are removed, it should be pointed out that all model success rates are well above 12.5%, the expectation for a random pick classification. Furthermore, for all models besides one exception, the correct model is clearly preferred over the others. The exception is *symmA*, which by raw numbers has more *f6* images associated to it then from its correct underlying class . However, also this trend is remedied once one considers the average similarity per class.

Physically interesting results can be seen in table A2, which shows the model classification success rates as a function of redshift. Two results catch the eye. Firstly, *symmA* is not classified at all at $z = 1$, which relates to the fact that this specific gravity model has no signature at this high redshift. Other degeneracies fully dominate in this case. Secondly, *dgp05* has a particularly high classification success rate at $z = 0.5$. This effect is not quite clear to us and we have to postpone its explanation to future studies. Finally, as observed earlier, classification success rate usually go down as a function of redshift, at it was already discussed in section 4.

Some tendencies can be understood from looking at other commonly studied observables. For example the similarities we see between *f6* and *symmA* is not that surprising given that these two models are also close when considering classical clustering observables. For example the matter power spectrum and halo mass function at $z = 0$ is very close in these two models.

To show a comparison of models which have all a similar value of σ_8 (compare section 4.3), we show in table A3 the classification matrices for training and test sets that contain *f5*, *symmB*, *dgp12* and a *lcdm* model that produces $\sigma_8 = 0.85$. The results agree well qualitatively with what has been seen earlier, but shift some of the degeneracies.

APPENDIX B: FULL FISHER SCORE STATISTICS FOR FEATURE FAMILIES

It would be excessive to list the fisher discriminant score for each of 2919 features in this image characterisation exercise. But to show at least the complete feature vector that was used in the analysis leading to section 4.1 we complete table 12 with table B1 showing the remaining features with their Fisher scores in the total feature vector with 73 elements.

Table B2 completes table 13 and now lists all 124 image feature classes that contain the 2919 image features, together with basis statistics on the Fisher scores within the class.

Table A1. Classification success and similarity matrix

	lcdm	f5	f6	symmA	symmB	dgp05	dgp12	dgp56	success rate
	73	20	33	30	33	4	5	2	
lcdm	0.1876	0.1364	0.1578	0.1638	0.1378	0.0771	0.0708	0.0686	0.37
	1.00	0.73	0.84	0.87	0.73	0.41	0.38	0.37	
	18	83	17	14	58	5	3	2	
f5	0.1452	0.1945	0.1522	0.1512	0.1639	0.0664	0.0639	0.0626	0.42
	0.75	1.00	0.78	0.78	0.84	0.34	0.33	0.32	
	32	19	52	39	42	5	8	3	
f6	0.1579	0.1464	0.1650	0.1769	0.1528	0.0704	0.0664	0.0644	0.26
	0.89	0.82	1.00	0.95	0.88	0.45	0.41	0.40	
	33	23	48	45	40	3	4	4	
symmA	0.1579	0.1464	0.1650	0.1769	0.1528	0.0704	0.0664	0.0644	0.23
	0.89	0.83	0.93	1.00	0.86	0.40	0.38	0.36	
	16	21	24	15	117	3	4	0	
symmB	0.1457	0.1536	0.1532	0.1580	0.2000	0.0648	0.0641	0.0607	0.59
	0.73	0.77	0.77	0.79	1.00	0.32	0.32	0.30	
	24	9	18	10	14	66	36	23	
dgp05	0.0876	0.0769	0.0875	0.0848	0.0794	0.2164	0.1908	0.1766	0.33
	0.40	0.36	0.40	0.39	0.37	1.00	0.88	0.82	
	24	9	16	15	17	31	56	32	
dgp12	0.0891	0.0772	0.0899	0.0880	0.0815	0.1839	0.2034	0.1871	0.28
	0.44	0.38	0.44	0.43	0.40	0.90	1.00	0.92	
	22	5	23	8	12	28	34	68	
dgp56	0.0827	0.0711	0.0845	0.0785	0.0720	0.1813	0.1953	0.2347	0.34
	0.35	0.30	0.36	0.33	0.31	0.77	0.83	1.00	
all									0.35

Table A2. Classification success rate at different redshifts

z	0.0	0.5	1.0
lcdm	0.37	0.27	0.29
f5	0.42	0.47	0.39
f6	0.26	0.24	0.26
symmA	0.23	0.25	**0.0**
symmB	0.59	0.40	0.31
dgp05	0.33	**0.82**	0.24
dgp12	0.28	0.29	0.20
dgp56	0.34	0.47	0.31
total	0.35	0.40	0.25

Table A3. Classification matrices with the *dgp12* model and with a $lcdm_{0.85}$ model with $\sigma_8 = 0.85$.

	$lcdm_{0.85}$	f5	symmB	dgp12	success rate
	111	15	36	38	
$lcdm_{0.85}$	0.4694	0.1115	0.1252	0.2939	0.56
	1.00	0.24	0.27	0.63	
	2	112	80	6	
f5	0.1003	0.4063	0.3532	0.1402	0.56
	0.25	1.00	0.87	0.35	
	2	45	147	6	
symmB	0.1032	0.3353	0.4209	0.1406	0.74
	0.25	0.80	1.00	0.33	
	32	30	47	91	
dgp12	0.2807	0.1763	0.1846	0.3583	0.46
	0.78	0.49	0.52	1.00	
all					0.58

Table B1. Table 12 continued

Rank	Name	Weight
26	Tamura Textures (Chebyshev ()) [5]	0.3426
27	Zernike Coefficients (Fourier (Wavelet ())) [71]	0.3379
28	Haralick Textures () [14]	0.3273
29	Haralick Textures (Edge ()) [14]	0.3211
30	Haralick Textures (Wavelet ()) [4]	0.3202
31	Zernike Coefficients (Fourier (Wavelet ())) [65]	0.3156
32	Haralick Textures (Edge ()) [11]	0.3154
33	Multiscale Histograms (Chebyshev (Fourier ())) [22]	0.3119
34	Zernike Coefficients (Fourier (Wavelet ())) [10]	0.3094
35	Haralick Textures () [10]	0.306
36	Haralick Textures () [8]	0.3028
37	Haralick Textures () [18]	0.2981
38	Zernike Coefficients (Fourier ()) [62]	0.292
39	Multiscale Histograms (Chebyshev (Fourier ())) [21]	0.2874
40	Zernike Coefficients (Fourier ()) [48]	0.287
41	Zernike Coefficients (Fourier ()) [39]	0.2864
42	Haralick Textures (Wavelet (Edge ())) [4]	0.2854
43	Haralick Textures (Edge ()) [9]	0.2821
44	Zernike Coefficients (Fourier ()) [24]	0.2801
45	Haralick Textures (Edge ()) [8]	0.2789
46	Zernike Coefficients (Fourier ()) [34]	0.2762
47	Haralick Textures () [22]	0.2688
48	Haralick Textures (Fourier (Chebyshev ())) [18]	0.2683
49	Zernike Coefficients (Fourier ()) [57]	0.265
50	Zernike Coefficients (Fourier (Wavelet ())) [37]	0.2646
51	Zernike Coefficients (Fourier ()) [8]	0.2602
52	Gini Coefficient (Wavelet (Fourier ())) [0]	0.259
53	Haralick Textures () [20]	0.258
54	Haralick Textures (Fourier (Chebyshev ())) [12]	0.2578
55	Zernike Coefficients (Fourier (Wavelet ())) [34]	0.2557
56	Zernike Coefficients (Fourier (Wavelet ())) [24]	0.2533
57	Zernike Coefficients (Fourier ()) [45]	0.2526
58	Haralick Textures (Wavelet (Edge ())) [12]	0.2517
59	Haralick Textures () [19]	0.2517
60	Zernike Coefficients (Fourier (Wavelet ())) [62]	0.2514
61	Haralick Textures () [1]	0.2502
62	Haralick Textures (Chebyshev (Wavelet ())) [14]	0.2495
63	Gini Coefficient (Chebyshev (Wavelet ())) [0]	0.2485
64	Zernike Coefficients (Fourier (Wavelet ())) [48]	0.2463
65	Zernike Coefficients (Fourier (Wavelet ())) [8]	0.2461
66	Haralick Textures (Chebyshev ()) [8]	0.2406
67	Haralick Textures (Chebyshev (Wavelet ())) [10]	0.2401
68	Gini Coefficient (Edge ()) [0]	0.2361
69	Tamura Textures (Fourier ()) [5]	0.2337
70	Zernike Coefficients (Fourier (Wavelet ())) [46]	0.2327
71	Haralick Textures (Chebyshev ()) [26]	0.2326
72	Haralick Textures (Chebyshev ()) [2]	0.2308
73	Zernike Coefficients (Fourier (Wavelet ())) [31]	0.2299

Table B2. Table 13 continued

Rank	Name	Min	Max	Mean	Std. dev.	#
16	Radon Coefficients (Fourier ())	0.002796	0.1572	0.07966	0.05538	12
17	Haralick Textures (Chebyshev (Wavelet ()))	0.001528	0.2495	0.07426	0.08442	28
18	Haralick Textures (Fourier (Chebyshev ()))	0.001438	0.2683	0.07406	0.08039	28
19	Haralick Textures (Chebyshev ())	0.0002557	0.2406	0.07387	0.08594	28
20	Tamura Textures (Wavelet (Fourier ()))	0.008465	0.1561	0.07253	0.05688	6
21	Radon Coefficients (Wavelet (Fourier ()))	0.0006693	0.1932	0.0725	0.06093	12
22	Radon Coefficients (Fourier (Edge ()))	0.01066	0.1438	0.06954	0.04716	12
23	Fractal Features (Fourier (Edge ()))	0.04659	0.07981	0.06902	0.01171	20
24	Radon Coefficients (Fourier (Wavelet ()))	0.005142	0.1236	0.06871	0.04525	12
25	Pixel Intensity Statistics (Chebyshev ())	0.003185	0.1221	0.06781	0.05914	5
26	Tamura Textures (Wavelet (Edge ()))	0.002864	0.1124	0.06567	0.04648	6
27	Gini Coefficient (Wavelet ())	0.06536	0.06536	0.06536	–	1
28	Multiscale Histograms (Chebyshev (Fourier ()))	2.761e-05	0.3119	0.06379	0.09224	24
29	Tamura Textures (Fourier (Wavelet ()))	0.0	0.1739	0.06093	0.08504	6
30	Haralick Textures (Wavelet (Fourier ()))	0.001761	0.2073	0.05406	0.05679	28
31	Edge Features ()	0.0	0.1329	0.05388	0.04086	28
32	Pixel Intensity Statistics (Edge ())	0.0	0.2079	0.05386	0.08666	5
33	Pixel Intensity Statistics (Fourier (Chebyshev ()))	0.001726	0.1099	0.05127	0.04356	5
34	Zernike Coefficients (Fourier (Edge ()))	0.002065	0.1793	0.0508	0.04392	72
35	Pixel Intensity Statistics (Wavelet (Edge ()))	0.004129	0.1825	0.05042	0.07427	5
36	Haralick Textures (Chebyshev (Fourier ()))	0.0007449	0.1662	0.04898	0.05044	28
37	Pixel Intensity Statistics (Wavelet (Fourier ()))	0.009989	0.102	0.04794	0.03919	5
38	Pixel Intensity Statistics (Chebyshev (Fourier ()))	0.0319	0.06915	0.0473	0.01647	5
39	Fractal Features (Chebyshev (Fourier ()))	0.03638	0.05426	0.04568	0.005604	20
40	Fractal Features (Fourier (Chebyshev ()))	0.004817	0.0715	0.04502	0.02301	20
41	Haralick Textures (Wavelet ())	0.0005351	0.3202	0.04411	0.07016	28
42	Radon Coefficients (Chebyshev (Wavelet ()))	0.0004625	0.09457	0.04259	0.04378	12
43	Tamura Textures (Wavelet ())	0.002278	0.07851	0.04257	0.02963	6
44	Pixel Intensity Statistics (Chebyshev (Wavelet ()))	0.0005042	0.0785	0.04232	0.03863	5
45	Chebyshev Coefficients (Fourier ())	0.0001485	0.1802	0.04178	0.05048	32
46	Gini Coefficient (Fourier (Chebyshev ()))	0.04167	0.04167	0.04167	–	1
47	Haralick Textures (Wavelet (Edge ()))	6.751e-05	0.2854	0.04003	0.06803	28
48	Pixel Intensity Statistics (Fourier (Edge ()))	0.006477	0.07215	0.0386	0.03069	5
49	Chebyshev-Fourier Coefficients ()	0.00278	0.08729	0.03718	0.02502	32
50	Tamura Textures (Chebyshev (Fourier ()))	0.0	0.1086	0.03678	0.03851	6
51	Haralick Textures (Fourier (Wavelet ()))	0.003524	0.09116	0.03545	0.02924	28
52	Pixel Intensity Statistics (Fourier ())	0.006817	0.09874	0.03542	0.04006	5
53	Gabor Textures ()	0.005368	0.05098	0.03541	0.01482	7
54	Multiscale Histograms ()	0.0	0.06035	0.03538	0.01127	24
55	Pixel Intensity Statistics (Fourier (Wavelet ()))	0.006814	0.08721	0.03479	0.03649	5
56	Chebyshev Coefficients (Fourier (Wavelet ()))	7.333e-05	0.1373	0.03357	0.0354	32
57	Fractal Features (Wavelet (Edge ()))	0.01615	0.08687	0.03158	0.01775	20
58	Haralick Textures (Fourier ())	0.005475	0.08353	0.03136	0.02556	28
59	Radon Coefficients (Chebyshev ())	0.001171	0.1105	0.03025	0.03528	12
60	Zernike Coefficients (Wavelet ())	0.001553	0.04834	0.02949	0.01235	72
61	Tamura Textures (Fourier (Chebyshev ()))	0.0	0.05934	0.0291	0.02429	6
62	Tamura Textures (Fourier (Edge ()))	0.0	0.07595	0.02848	0.03662	6
63	Comb Moments (Edge ())	0.003435	0.1302	0.02723	0.02894	48
64	Multiscale Histograms (Chebyshev ())	0.001093	0.08173	0.02662	0.02668	24
65	Multiscale Histograms (Chebyshev (Wavelet ()))	0.00025	0.07265	0.02647	0.01974	24
66	Comb Moments (Fourier (Wavelet ()))	0.0	0.1406	0.02633	0.039	48
67	Multiscale Histograms (Fourier (Wavelet ()))	0.0	0.06156	0.02345	0.01592	24
68	Fractal Features (Edge ())	0.009364	0.06868	0.02333	0.01531	20
69	Zernike Coefficients ()	0.0004158	0.04813	0.02311	0.01263	72
70	Gini Coefficient (Fourier (Edge ()))	0.02216	0.02216	0.02216	–	1

Table B3. Table 13 continued

Rank	Name	Min	Max	Mean	Std. dev.	#
71	Tamura Textures (Edge ())	0.0006446	0.05768	0.02165	0.02062	6
72	Pixel Intensity Statistics (Wavelet ())	0.006814	0.05038	0.02161	0.01707	5
73	Chebyshev Coefficients ()	0.0008869	0.0634	0.02085	0.01768	32
74	Radon Coefficients (Chebyshev (Fourier ()))	0.0008424	0.05026	0.02029	0.01923	12
75	Multiscale Histograms (Wavelet (Fourier ()))	0.0002537	0.08552	0.02007	0.02656	24
76	Comb Moments ()	0.000969	0.0631	0.01988	0.01873	48
77	Multiscale Histograms (Edge ())	0.0	0.04351	0.01965	0.01608	24
78	Comb Moments (Chebyshev (Wavelet ()))	2.809e-05	0.1204	0.01957	0.02995	48
79	Chebyshev Coefficients (Fourier (Edge ()))	7.002e-05	0.1566	0.01886	0.03433	32
80	Comb Moments (Fourier (Edge ()))	0.0	0.06139	0.0188	0.01759	48
81	Comb Moments (Chebyshev (Fourier ()))	0.0001936	0.07083	0.01856	0.01756	48
82	Multiscale Histograms (Fourier ())	0.0	0.03878	0.01818	0.01083	24
83	Haralick Textures (Fourier (Edge ()))	0.0011	0.0498	0.01772	0.01365	28
84	Chebyshev Coefficients (Edge ())	2.742e-05	0.05939	0.01769	0.01718	32
85	Pixel Intensity Statistics ()	0.0	0.05479	0.01738	0.02333	5
86	Multiscale Histograms (Wavelet ())	0.0	0.03136	0.01676	0.011	24
87	Tamura Textures ()	0.003395	0.03918	0.01556	0.01372	6
88	Comb Moments (Fourier ())	0.0	0.08391	0.01543	0.02206	48
89	Chebyshev-Fourier Coefficients (Edge ())	0.002144	0.03468	0.0149	0.008587	32
90	Otsu Object Features ()	2.923e-05	0.07006	0.01471	0.01543	34
91	Multiscale Histograms (Fourier (Edge ()))	0.0	0.03262	0.01357	0.01125	24
92	Comb Moments (Fourier (Chebyshev ()))	0.0002988	0.06055	0.01302	0.01223	48
93	Comb Moments (Wavelet (Edge ()))	0.0005041	0.08158	0.01256	0.01493	48
94	Comb Moments (Wavelet ())	0.000263	0.05191	0.01214	0.01211	48
95	Chebyshev Coefficients (Wavelet ())	0.0006492	0.03937	0.01193	0.009978	32
96	Fractal Features (Fourier (Wavelet ()))	0.006856	0.03016	0.01164	0.005313	20
97	Fractal Features ()	0.003373	0.02253	0.01111	0.006299	20
98	Comb Moments (Wavelet (Fourier ()))	0.0	0.06691	0.01077	0.01314	48
99	Zernike Coefficients (Edge ())	0.0009013	0.03445	0.01066	0.006926	72
100	Fractal Features (Chebyshev (Wavelet ()))	0.004038	0.02374	0.01037	0.004437	20
101	Fractal Features (Wavelet ())	0.001692	0.02562	0.01001	0.006148	20
102	Fractal Features (Wavelet (Fourier ()))	0.00538	0.0117	0.009496	0.002046	20
103	Zernike Coefficients (Wavelet (Edge ()))	0.0008492	0.02444	0.009097	0.005664	72
104	Gini Coefficient (Chebyshev (Fourier ()))	0.00895	0.00895	0.00895	–	1
105	Radon Coefficients (Fourier (Chebyshev ()))	0.002551	0.0178	0.008616	0.005505	12
106	Chebyshev Coefficients (Chebyshev ())	0.0002586	0.01941	0.008269	0.006129	32
107	Multiscale Histograms (Wavelet (Edge ()))	0.0001181	0.02458	0.008168	0.009393	24
108	Fractal Features (Fourier ())	0.005944	0.02786	0.008154	0.004887	20
109	Multiscale Histograms (Fourier (Chebyshev ()))	0.0	0.01491	0.00687	0.004261	24
110	Chebyshev-Fourier Coefficients (Fourier (Wavelet ()))	0.0002502	0.05058	0.006706	0.01031	32
111	Chebyshev Coefficients (Wavelet (Edge ()))	0.0005705	0.02501	0.006566	0.005566	32
112	Radon Coefficients (Wavelet (Edge ()))	0.001386	0.0106	0.006331	0.002556	12
113	Fractal Features (Chebyshev ())	0.004724	0.007529	0.006328	0.0009132	20
114	Chebyshev-Fourier Coefficients (Fourier (Edge ()))	7.099e-06	0.02821	0.005849	0.006827	32
115	Chebyshev-Fourier Coefficients (Wavelet (Edge ()))	0.0004125	0.01329	0.005519	0.003601	32
116	Radon Coefficients (Edge ())	0.001494	0.007173	0.005212	0.001605	12
117	Comb Moments (Chebyshev ())	4.192e-05	0.04941	0.005006	0.01017	48
118	Radon Coefficients (Wavelet ())	0.0006035	0.008731	0.00455	0.002576	12
119	Radon Coefficients ()	0.0003923	0.01101	0.004297	0.003788	12
120	Chebyshev-Fourier Coefficients (Fourier ())	3.925e-05	0.01895	0.003691	0.004664	32
121	Chebyshev-Fourier Coefficients (Wavelet ())	0.0002408	0.01188	0.003343	0.002948	32
122	Chebyshev-Fourier Coefficients (Chebyshev ())	0.000334	0.004244	0.001783	0.001106	32
123	Inverse-Otsu Object Features ()	0.0	0.01933	0.001535	0.004358	34
124	Zernike Coefficients (Chebyshev ())	0.001355	0.001563	0.001451	5.332e-05	72